

Modelling and inference for the movement of interacting animals

Jordan E. Milner¹  | Paul G. Blackwell¹  | Mu Niu² 

¹School of Mathematics & Statistics,
University of Sheffield, Sheffield, UK

²School of Computing, Electronics and
Mathematics, University of Plymouth,
Plymouth, UK

Correspondence

Jordan E. Milner
Email: jemilner1@sheffield.ac.uk

Present address

Mu Niu, School of Mathematics and
Statistics, University of Glasgow, Glasgow,
GL 8QQ, UK

Funding information

Engineering and Physical Sciences Research
Council

Handling Editor: Timothée Poisot

Abstract

1. Statistical modelling of animal movement data is a rapidly growing area of research. Typically though, these models have been developed for analysing the tracks of individual animals and we lose sight of the impact animals have on each other with regards to their movement behaviours. We aim to develop a model with a flexible social framework that allows us to capture that information.
2. Our approach is based on the concept of social hierarchies, and this is embedded in a multivariate diffusion process which models the movement of a group of animals. The possibility of switching between behavioural states facilitates dynamic social behaviours and we augment the observed data with sampled state switching times in order to model the animals' behaviour naturally in continuous time. In addition, this enables us to carry out exact inference in a Bayesian setting with the benefits of being able to handle regular, irregular and missing data. All movement and behaviour parameters are estimated with Markov chain Monte Carlo methods.
3. We examine the capability of our model with simulated data before fitting it to GPS locations of five wild olive baboons *Papio anubis*. The results enable us to identify which animals are influencing the movement of others and when, which provides both a dynamic and long-term static insight into the group's social behaviours.
4. Our model offers a flexible method in continuous time with which to model the network of social interactions within animal movement. Doing so avoids the limitations caused by a discrete-time approach and it allows us to capture rich information with regards to a group's social structure, leading to constructive applications in conservation and management decisions. However, currently it is a computationally expensive task to fit the model to data, which in turns limits extending the model to more fruitful but complex cases such as heterogeneity in space or individual characteristics. Furthermore, our social hierarchy approach assumes all relevant animals are tracked and that any interactions have some ordering, both of which narrow the scope within which this approach is appropriate.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Methods in Ecology and Evolution* published by John Wiley & Sons Ltd on behalf of British Ecological Society

KEYWORDS

Bayesian statistics, behaviour state switching, collective movement modelling, continuous time, diffusion process, MCMC, social animals, wild baboons

1 | INTRODUCTION

The research area of statistically modelling animal movement has rapidly expanded in recent years. This has in part been driven by the increase in availability of movement data (e.g. from GPS tags) coupled with the potential insight, from habitat preference to monitoring the impacts of a changing environment, that can be gained from analysing it (Kays, Crofoot, Jetz, & Wikelski, 2015).

This has led to a broad range of methods to be developed. For instance, the hidden Markov model approach utilises the computational efficiency of discrete time (Langrock et al., 2012) whilst the continuous-time correlated random walk formulation by Johnson, London, Lea, and Durban (2008) offers more flexibility with regards to irregular data intervals. Dunn and Gipson (1977) introduced modelling animal movement as a diffusion process, a concept that has been built up on by Blackwell (1997, 2003), Harris and Blackwell (2013) and Blackwell, Niu, Lambert, and LaPoint (2016) to account for multiple movement behaviours.

Typically, these methods have been developed for analysis of individual animals and so they fail to account for the impacts that social animals have on each other's movement behaviours. With increasing ability to obtain simultaneous tracking data from multiple animals within a group (Westley, Berdahl, Torney, & Biro, 2018), we now have the opportunity to build these social interactions into our models. Indeed, recent work has begun to explore that possibility (Langrock et al., 2014; Niu, Blackwell, & Skarin, 2016) through treating the group as a collective during the movement.

However, being able to drill further into the social behaviour of a group, such as identifying animals with high levels of influence, will provide us with richer information on their social structures, with useful applications in conservation efforts (King, Fehlmann, Biro, Ward, & Fürtbauer, 2018; Westley et al., 2018). Outside of statistical modelling, approaches have been taken to extract this desired information from within the data, such as Strandburg-Peshkin, Farine, Couzin, and Crofoot (2015) looking for moments where collective movement decisions have occurred and how the scenario developed.

In this paper, we develop a social framework in our movement model to capture those intra-group interactions. This is inspired by orderly social hierarchies, which are shown to be prevalent across a broad array of taxa (McDonald & Shizuka, 2013). The need to keep these hierarchies dynamic (Chase & Lindquist, 2016) also works in conjunction with more fluid arrangements such as fission-fusion dynamics (Ramos-Fernández & Morales, 2014) meaning we can account for an extensive range of social constructs.

Furthermore, we formulate this in continuous time. In doing so, the interactive behaviours we define are not married to the temporal scale of the observations, we avoid the approximations caused from

analysing discrete data and irregular or missing data are not problematic (Blackwell et al., 2016; Harris & Blackwell, 2013).

During this paper, we first introduce the components of our movement model: the social framework which is embedded within a multivariate Ornstein-Uhlenbeck process and continuous-time behavioural state switching. We then outline the algorithm developed to fit the model to data using Markov chain Monte Carlo (MCMC) methods before showcasing key features of the results we can obtain using simulated data and baboon GPS locations. Finally, we discuss the wider applications of this work and future directions.

2 | MODELLING SOCIAL ANIMAL MOVEMENT

2.1 | Influence hierarchies

To capture the social behaviour of a group's movement, our assumption is: a period of direct interaction between two animals can be characterised by the movement of one of those animals being attracted to the other. This is an assumption shared with other areas of animal movement literature (Long, Nelson, Webb, & Gee, 2014).

For ease of reference, we will refer to the roles in this dyadic relationship as 'dominant' and 'subordinate' (see Social definitions) but, as with all behaviour modelling, we need to be careful not to over-interpret the behaviour labelling. That is, 'subordinate' or 'attracted to' have certain connotations but the movement behaviour just broadly translates to the movement of an animal being influenced by the other in some sense. Similar considerations also need to be made for our subgroup-level labels of 'leader' and 'follower'.

This dyadic concept can be naturally extended for larger social groups to give rise to social hierarchies (see Figure 1a). In order to keep this framework tractable and easy to interpret, we are restricting these hierarchies to essentially a thinned network that contain the edges representing the most causal interaction. That is, an animal can have at most one dominant but it can have multiple subordinates. This is as opposed to a tournament-style network where there is a some degree of relationship between every pair of nodes, which is a rare occurrence in nature (McDonald & Shizuka, 2013). Thus, our resulting social structure will represent the most causal influence to explain a group's movement—hence 'influence hierarchies' (see Figure 1b). We restrict the possible hierarchies to avoid any cycles, to ensure that the pattern of relationships is meaningful and that the movement models will be well-defined. The hierarchical structure is therefore what is often known in statistical contexts as a Directed Acyclic Graph.




FIGURE 1 In (a), whilst C is subordinate to B, we see B is in turn subordinate to A which enables us to learn how the influence in movement is cascaded through the group. In (b), we have a transitive triad where A dominates B, B dominates C and A dominates C. We only capture the influence that best describes C's movement and consequently the structure we estimate will either be A dominates B, B dominates C or A dominates B, A dominates C. In all of the above hierarchies, A is the leader whilst B and C are followers

Social definitions

Dominant: with respect to an edge in the hierarchy between animals i and j where j is attracted to i , i is dominant towards j .

Follower: an animal in a subgroup that is not the leader.

Group: all animals in the data set.

Independent: an animal is independent if it is neither a dominant or a subordinate.

Leader: the focal point of a subgroup. That is, this animal is a dominant to at least one animal, but a subordinate to none. Animal i 's leader is the leading animal of the subgroup i is in.

Subgroup: all animals in a single hierarchy structure, i.e. a component of the graph formed as in Figure 1. Independent animals are their own subgroup.

Subordinate: with respect to an edge in the hierarchy between animals i and j where j is attracted to i , j is subordinate to i .

2.2 | Multivariate Ornstein-Uhlenbeck process

To model the movement of the animals, we use a diffusion process including a linear attraction term to represent the attraction-based interaction we have assumed. Then the movement of an individual animal i that is attracted to animal j can be described by the following stochastic differential equation (SDE):

$$dA_{i_t}^y = -\alpha (A_{i_t}^y - A_{j_t}^y) dt + \sigma dW_{i_t}^y, \quad (1)$$

where $A_{i_t}^y$ is the location of animal i at time t in the y coordinate; α is the rate of attraction towards $A_{j_t}^y$ where $i \neq j$; σ is the coefficient of 'noise', the component of movement modelled not in terms of social

interaction but as Brownian motion $W_{i_t}^y$. Therefore, Equation 1 has two components: the noise term, which is a continuous-time analogue of a random walk, and the attraction term, which captures any persistence in the movement towards another animal. For leading and independent (i.e. non-subordinate) animals, this reduces to Brownian motion (BM) as they have no attraction term:

$$dA_{i_t}^y = \rho dW_{i_t}^y, \quad (2)$$

where ρ is a distinct noise parameter. The x and y coordinates are treated as independent [see Blackwell (1997) for justification] and there are corresponding equations for the x axis.

Because of their linearity, these univariate SDEs can be combined into a multivariate Ornstein-Uhlenbeck (OU) process to model the group's movement jointly as detailed by Niu et al. (2016). For n animals, say, in the y axis:

$$d\mathbf{G}_t = \mathbf{F}_t (\mathbf{G}_t - \Theta_t) dt + \Sigma d\mathbf{B}_t, \quad (3)$$

where

$$\mathbf{G}_t = \begin{pmatrix} A_{1_t}^y \\ \vdots \\ A_{n_t}^y \end{pmatrix}, \Theta_t = \begin{pmatrix} L_{1_t}^y \\ \vdots \\ L_{n_t}^y \end{pmatrix}, \mathbf{B}_t = \begin{pmatrix} W_{1_t}^y \\ \vdots \\ W_{n_t}^y \end{pmatrix},$$

and $L_{i_t}^y$ is the location of animal i 's leader at time t . \mathbf{F}_t and Σ are $(n \times n)$ -matrices where

$$\mathbf{F}_{t_{ij}} = \begin{cases} -\alpha, & i = j \text{ and } i \text{ is a subordinate} \\ \alpha, & i \neq j \text{ and } i \text{ is subordinate to } j \\ 0, & \text{otherwise,} \end{cases}$$

and

$$\Sigma_{i_{ij}} = \begin{cases} \sigma, & i = j \text{ and } i \text{ is a subordinate} \\ \rho, & i = j \text{ and } i \text{ is a leader or independent} \\ 0, & \text{otherwise.} \end{cases}$$

\mathbf{G}_t is a vector of locations for all animals at time t and matrix \mathbf{F}_t is the attraction matrix for the group which indicates the interactions within the hierarchy at time t , i.e. who is subordinate to whom. Θ_t is a vector which contains the location of each animal's leader at time t and matrix Σ contains the coefficient of noise for each animal.

The solution to the multivariate SDE has a closed form—a multivariate normal distribution (Niu et al., 2016):

$$\mathbf{G}_t | \mathbf{G}_0 \sim \text{MVN}(\boldsymbol{\mu}(\mathbf{G}_0, \mathbf{F}_0, t), \Xi(\mathbf{F}_0, t)), \quad (4)$$

where \mathbf{G}_0 , \mathbf{F}_0 and Θ_0 correspond to the animals' locations, the attraction matrix and the leaders' locations at time 0 respectively. The OU

process is Markovian with the animals' locations at time t conditional on their previous locations at time 0.

The expected value of this distribution is given by:

$$E[\mathbf{G}_t | \mathbf{G}_0] = \boldsymbol{\mu}(\mathbf{G}_0, \mathbf{F}_0, t) = e^{\mathbf{F}_0 t} (\mathbf{G}_0 - \Theta_0) + \Theta_0, \quad (5)$$

and $\text{Var}[\mathbf{G}_t | \mathbf{G}_0]$ is given by $\Xi(\mathbf{F}_0, t)$ which consists of the following five expressions (to ease notation, we simplify $\Xi(\mathbf{F}_0, t)$ to Ξ and \mathbf{F}_0 to \mathbf{F}):

$$\Xi_{ij} = \Xi_{ji} = \begin{cases} \rho^2 t, & (a) \\ \Xi_{\text{dom}(i)j} - \frac{\rho^2 e^{\mathbf{F}t}_{ij}}{\alpha}, & (b) \\ \frac{\Xi_{\text{dom}(i)j} + \Xi_{i,\text{dom}(i)}}{2} - \frac{\rho^2 e^{\mathbf{F}t}_{ij} e^{\mathbf{F}t}_{j,l}}{2\alpha} - \frac{\sigma^2 (e^{\mathbf{F}t}_{i,-l} \cdot e^{\mathbf{F}t}_{j,-l})}{2\alpha}, & (c) \\ \Xi_{i,\text{dom}(i)} - \frac{\rho^2 e^{\mathbf{F}t}_{i,l} e^{\mathbf{F}t}_{l,j}}{2\alpha} + \frac{\sigma^2 (1 - e^{\mathbf{F}t}_{i,-l} \cdot e^{\mathbf{F}t}_{l,-l})}{2\alpha}, & (d) \\ 0, & (e) \end{cases} \quad (6)$$

where $\text{dom}(i)$ is the dominant of animal i ; l is the leader of the subgroup both i and j are in; $-l$ indicates all animals except l . The scenarios of (a) to (e) are as follows:

- $i = j$ and i is a leading or independent animal.
- $i \neq j$ and j is the leader of i 's subgroup.
- $i \neq j$, i and j are in the same subgroup but neither are the leader.
- $i = j$ and i is a subordinate.
- $i \neq j$ and i and j are in different subgroups.

In practice, the order of the computation of $\Xi(\mathbf{F}_0, t)$ is important as some expressions rely on other values within $\Xi(\mathbf{F}_0, t)$. See Appendix A for the algorithm (all appendices are provided in the online supporting information).

2.3 | Behaviour states

The movement model we have described so far relies on some knowledge of which animal is subordinate to which and when. However, we are unlikely to know either of those pieces of information and so we treat them as unknown. To estimate them, we incorporate behavioural state switching where the states correspond to the animal's social behaviour. As we are operating in continuous time, we do so with a continuous-time Markov chain with a discrete state space.

The state space we want to explore is the space of all possible hierarchies as defined in Section 2.1 and we are restricting each switching time to contain at most one behaviour change. That is, at most one animal can change behaviour at a given time. Ideally, the transition rates of our Markov chain would correspond to switching between hierarchies. However, even for a group as small as four animals there are in excess of 100 of these structures, meaning that approach is not practical. Whilst we would set the majority of the

transition rates to be 0 due to the above restriction, the transition matrix would be unwieldy and difficult to define and interpret.

We can, however, explore the same state space by defining our behaviour states to represent each individual animal's state, as opposed to the group's. That is, an individual can switch which animal they are attracted to (i.e. subordinate to) or switch to Brownian motion (i.e. leading or independent). The following generator matrix is then used for each animal:

$$\Lambda = \begin{matrix} & S_{A1} & \dots & S_{An} & \text{BM} \\ S_{A1} & \left(\begin{array}{cccc} \lambda_{1,1} & \dots & \dots & \lambda_{1,n+1} \\ \vdots & \ddots & & \vdots \\ S_{An} & \vdots & & \vdots \\ \text{BM} & \lambda_{n+1,1} & \dots & \dots & \lambda_{n+1,n+1} \end{array} \right) & & & \end{matrix}, \quad (7)$$

where n is the number of animals in the data and state S_{Ai} represents being subordinate to animal i . Here we use the parameterisation of a continuous-time Markov chain in which an animal stays in state u for a holding time, which is exponentially distributed with rate λ_u where λ_u is the rate of leaving state u ($\lambda_u = \sum_{u \neq v} \lambda_{uv}$). After this holding time, the animal switches to state v with probability λ_{uv}/λ_u . The diagonal elements of λ are therefore determined by the off-diagonal ones: $\lambda_{uu} = -\lambda_u$.

In the case of $n=4$ as in the above example, the number of transition parameters is a manageable 20 as well as being more intuitive. The penalty of using this individualistic approach as a proxy for switching hierarchies is akin to treating each animal homogeneously. That is, the transition rate λ_{uv} is the rate of switching from state u to v averaged over all animals in the group.

2.3.1 | Leading/independent extensions

The model described so far assumes leading or independent animals to be restricted to a single BM state and movement parameter ρ . To capture richer movement for these animals, various extensions to the model can be made.

Firstly, additional BM states with distinct noise parameters can be included to represent different 'speeds' of movement. In the results to be discussed, we use two BM states to loosely embody slow and fast movement. Secondly, we might assume that leading or independent animals are themselves attracted to some location, a resource or nesting site for example. It would then be natural to think of these animals as also moving under an OU process.

Both methods add movement parameter(s) and may add further row(s) and column(s) to the transition matrix, meaning some consideration is needed as to the costs and benefits of these extensions.

3 | INFERENCE

Markov chain Monte Carlo methods are used to infer both the behaviour and movement parameters. Each iteration of the MCMC

algorithm consists of two parts. Firstly, we sample the behaviour states of the animals in continuous time. This is done through simulating the state switches between the observed data points and the acceptance of these behaviour trajectories is evaluated by a Metropolis-Hastings ratio. The exact algorithm for this step is detailed in Section 3.1. Secondly, the parameters of the Ornstein-Uhlenbeck and Brownian motion processes are sampled in accordance with the latest behaviour sample by means of a Metropolis-Hastings random walk. Further details of this, including restrictions we place on the parameters, are in Section 3.2.

3.1 | Behaviour parameters

To estimate the behaviour parameters (both the states and Λ) we sample the animals' behaviour trajectories between the observed data points. Through this, we treat the behaviour switching in continuous time and account for our uncertainty in what the animal is doing at unobserved times, which enables us to undertake exact inference in a Bayesian setting. Our method to simulate these trajectories is a multi-animal variation of the 'kappa' method introduced by Blackwell et al. (2016) where the observed data is augmented with sampled state switching times. The current paper only discusses the spatially homogeneous case and we restrict only one animal to switch state at any given switching time. Furthermore, we only update the trajectory for one animal at a time to increase the acceptance rates of our proposals.

Let τ_o represent the observed data times for the interval $[t_a, t_b]$; τ_s represent the sampled switching times in the same interval for animal i ; $\tau = \tau_o \cup \tau_s$ so that $\tau = \{\tau_1 < \dots < \tau_p\}$ where p is the size of τ and $\tau_s = \cup_j \tau_{s_j}$ for all animals i . Finally, let β represent the behaviour states of all animals at times τ where β_j corresponds to the states at time τ_j for $j = 1, \dots, p$.

To re-sample a trajectory estimate for a given animal, say i , in the interval $[t_a, t_b]$, we discard the current sampled switching times for i , τ_{s_j} , and sample new ones, $\hat{\tau}_{s_j}$. The new potential switching times are produced from a homogeneous Poisson process over the interval (t_a, t_b) with rate λ_{\max} , where $\lambda_{\max} \geq \max(\lambda_u)$ for all states u . $\hat{\tau}_{s_j}$ is combined with τ_{s_j} , the remainder of τ_s , to produce a new set of switching times, $\hat{\tau}_{s_j}$, and augmented times, $\hat{\tau} = \tau_o \cup \hat{\tau}_{s_j}$.

We then simulate the behaviour states forward through $\hat{\tau} = \{\hat{\tau}_1 < \dots < \hat{\tau}_{\hat{p}}\}$, where \hat{p} is the size of $\hat{\tau}$, to obtain our new behaviour trajectory $\hat{\beta}$. We initialise $\hat{\beta}_1 = \beta_1$, after which there are three scenarios to account for:

- If $\hat{\tau}_j \in \tau_o$ for $j=2, \dots, \hat{p}$, the behaviour states of all animals at $\hat{\tau}_j$ are carried forward from $\hat{\beta}_{j-1}$.
- If $\hat{\tau}_j \in \tau_{s_{-j}}$ for $j = 2, \dots, \hat{p}-1$, we use our previously sampled states for all animals except i at these times whilst the behaviour state of i is carried forward from $\hat{\beta}_{j-1}$.
- If $\hat{\tau}_j \in \hat{\tau}_{s_j}$ for $j=2, \dots, \hat{p}-1$, the behaviour states of all animals except i at $\hat{\tau}_j$ are carried forward from $\hat{\beta}_{j-1}$. The probability of $\hat{\tau}_j$ being a switch for i is λ_u/λ_{\max} when i is in state u . If so, the new

state is v with probability λ_{uv}/λ_u , otherwise, the state of i at $\hat{\tau}_j$ is also carried forward from $\hat{\beta}_{j-1}$.

We initialise $\hat{\beta}_1 = \beta_1$ as we typically simulate a trajectory estimate over a short interval of the data and therefore our simulations must be consistent with the trajectories outside of that interval. For that reason, we also require $\hat{\beta}_{\hat{p}} = \beta_p$. Within the trajectory, there is an additional condition that sampled states must not create a cyclic hierarchy, which includes an animal not being able to be subordinate to itself. If these conditions aren't met, we reject the trajectory and return to our previously sampled switches for i .

If these conditions are met, we accept or reject the new trajectory with a Metropolis-Hastings (MH) step. Our simulations are proposed from the current estimate of Λ and so the MH ratio simplifies to a ratio of likelihoods of the observed movement through the new proposed state switches against our previous estimate:

$$\prod_{k=2}^{|\tau_o|} \frac{f(\mathbf{g}_{\tau_{o_k}} | \mathbf{g}_{\tau_{o_{k-1}}}, \hat{\beta}_{[\tau_{o_{k-1}}, \tau_{o_k}]}, \hat{\tau}_{[\tau_{o_{k-1}}, \tau_{o_k}]})}{f(\mathbf{g}_{\tau_{o_k}} | \mathbf{g}_{\tau_{o_{k-1}}}, \beta_{[\tau_{o_{k-1}}, \tau_{o_k}]}, \tau_{[\tau_{o_{k-1}}, \tau_{o_k}]})}, \quad (8)$$

where $\mathbf{g}_{\tau_{o_k}}$ are the locations of the animals in a particular axis at time τ_{o_k} ; $\hat{\beta}_{[\tau_{o_{k-1}}, \tau_{o_k}]}$ are the newly sampled states throughout the interval $[\tau_{o_{k-1}}, \tau_{o_k}]$ at times $\hat{\tau}_{[\tau_{o_{k-1}}, \tau_{o_k}]}$; $\beta_{[\tau_{o_{k-1}}, \tau_{o_k}]}$ are the previous state estimations in the same interval at times $\tau_{[\tau_{o_{k-1}}, \tau_{o_k}]}$. Blackwell (2003) details how the movement likelihood terms are calculated through behaviour switches between two observations and more detail can be found in Appendix B.

We use the conjugate Dirichlet prior for the multinomial likelihood of the transition rates to obtain their full conditional distributions and resample them using Gibbs sampling at each iteration of the MCMC algorithm.

3.1.1 | Partial observations

Partial observations, where we only have data on some of the animals we are tracking, are a potential obstacle when analysing data from multiple animals: tracking equipment may not be fully synchronised or a GPS tag may not have been able to transmit some data for example. However, the above method is naturally adaptable to take into account the uncertainty of any missing or unsynchronised data using standard results for conditional multivariate normal distributions. See Appendix C for more information.

3.2 | Movement parameters

Treating the current values of the behaviours as fixed, we update the parameters of all Ornstein-Uhlenbeck and Brownian motion processes simultaneously through a Metropolis-Hastings random walk. We use independent, normally-distributed proposals for each parameter meaning the Metropolis-Hastings ratio is again reduced to Equation 8.

We place certain bounds on the parameters. In general, all parameters must be >0 and we have added the restriction that $\sigma \leq \rho$, the justification being that an animal influenced by another should have less variability in their movement than an animal following a BM process. In the case of extending the model to allow for multiple BM states (see Section 2.3), the above restriction is loosened to $\sigma \leq \rho_{\max}$ where $\rho_{\max} \geq \rho, \forall \rho$. This case also brings potential complications to the state labelling and so, in order to keep consistency, we define $\rho_1 < \dots < \rho_m$ for m BM states.

3.3 | Implementation

Both the simulation and inference methods were fully implemented in R (R Core Team, 2017). Aside from the `MCMCPACK` package (Martin, Quinn, & Park, 2011), which is used for the Dirichlet distribution, the code to perform the inference is original. Whilst the inference code is not yet formally wrapped up as an R package, it is available on GitHub (see Data Availability Statement) along with a brief readme file that instructs on how to use the code in R. This includes guidance on what format the data is required to be in, what tuning parameters need consideration and altering the number of states which model leading and independent movement. There is also information for reproducing the analysis in this paper.

4 | RESULTS

4.1 | Simulated data

We analysed simulated data to show the theoretical capabilities of both our model and inference approach. Here, we analyse a single data set to examine the outputs of the model in detail, whilst in Section 5, we analyse 400 different simulations to provide insight into the robustness of the model in different scenarios.

The simulation consists of five dynamically interacting animals for 100 discrete-time movement steps, with each step two units of time. We randomly deleted 10% of the data (uniformly across all data) to provide us with the setting of having incomplete

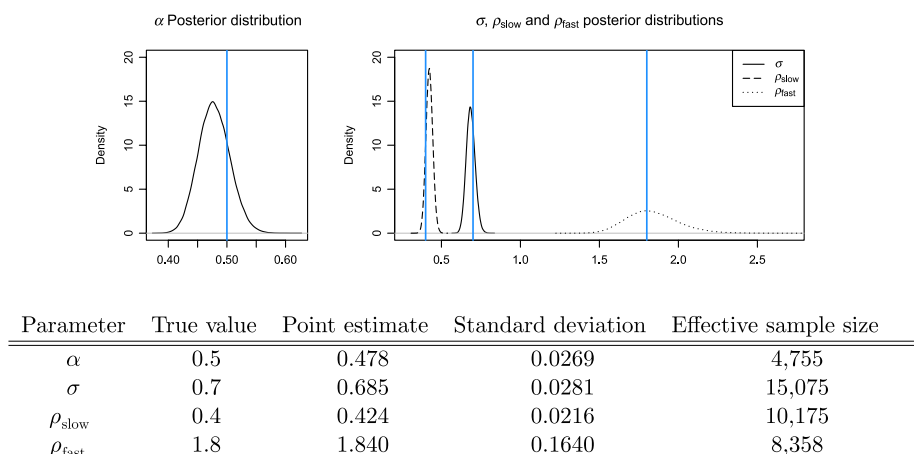
data. The ‘observed’ times were augmented with switching times through a Poisson process as detailed in Section 3.1; the movement of the group was simulated forwards through the augmented data set using a multivariate OU process as in Section 2.2; state switches were sampled using a continuous-time Markov chain as in Section 2.3.

We incorporated two BM states for leading/independent animals to allow for different speeds of movement, meaning we have seven behaviour states in total. Whilst we simulated the movement and behaviours of the animals in continuous time, the model is only fitted to the discrete-time ‘observed’ data to mirror typical real data. All inference runs (in this and subsequent sections) were performed on the University of Sheffield’s HPC ‘ShARC’ where each core runs at 2.4 GHz with 4.0 GB of RAM (a single core being used for each run).

To initialise the MCMC algorithm we randomised the behaviour parameters and started with over-dispersed movement parameter values. We set λ_{\max} as 0.2 as that was sufficiently high for the transition rates used for the simulation and we updated sections of an animal’s trajectory ranging from 3 to 12 observations long. We sampled a new trajectory 70 times per iteration of the MCMC in order for the behaviour state at each observation to be re-sampled on average. We ran the MCMC algorithm for 1.4 million iterations, of which 50,000 was burn-in, and we recorded every second iteration for the movement parameters and every 20th iteration for the behaviour parameters. We used an uninformative Dirichlet prior for the transition rates.

Figure 2 shows the movement parameter posteriors against the true values used. All are consistent with the true value and to check convergence the Gelman-Rubin diagnostic is used to assess the potential improvement from running more or longer chains. The multivariate potential scale reduction factor (PSFR) is 1, calculated from two separate MCMC runs using *coda* (Plummer, Best, Cowles, & Vines, 2006); this indicates that each chain is exploring the same posterior distribution, after burn-in. Details of the confidence limits of the PSFRs (for these and subsequent results) are in Appendix D. Figure 3 shows our behaviour state posteriors for animal 3, along with the true states from the simulation. The posterior estimates are broadly correct and confident, though quick, nuanced switches as in observations 67 and 68 can be smoothed

FIGURE 2 Top: posterior distributions for the four movement parameters for our simulated data. The blue vertical line indicates the true value used. Bottom: a summary of the movement parameter results. Point estimates and standard deviations are given to 3 s.f.; effective sample size is rounded down after being calculated using using *coda* (Plummer et al., 2006)






FIGURE 3 The state posterior distribution for animal 3 in the simulation data. There are seven states: the two Brownian motion speeds and five subordinate behaviours where state S_{A_i} indicates attraction to animal i . The area of each box represents the posterior probability of being in that state at that observation, from 0 to 1. The black line is true state in the simulated data

over. Note that since this output refers to animal 3, the probability of it being in state S_{A_3} (which would represent subordinate to itself) is necessarily zero.

Thus, our model and inference approach can provide insights into the dynamics of the social interactions within a group's movement when our social behaviour assumptions (Section 2.1) are fair.

4.2 | Baboon data

To test our approach 'in the field', we have taken a subset of wild olive baboon data that was originally analysed by Strandburg-Peshkin et al. (2015), which is available on Movebank (Crofoot, Kays, & Wikelski, 2015). The GPS data was collected at the Mpala Research Centre in Kenya for 26 baboons in a single troop. The data was recorded at a frequency of 1 Hz, for 12 hr a day (06:00–18:00) over 30 days. We took a subset of this data for five baboons (ID's 3, 4, 5, 11, 9) for 15 min (899 observations) to act as a test for the model. We chose this time period to contain some directional conflict as in Movie S2 in the supplementary materials of Strandburg-Peshkin et al. (2015). We converted the GPS coordinates to UTM zone 37N easting-northing using *sp* (Bivand, Pebesma, & Gomez-Rubio, 2013; Pebesma & Bivand, 2005).

Again, we included two BM states for leading or independent animals. In attempting to fit the model with only a single BM state in another analysis, the subordinate behaviour states translated to a pseudo-BM state (that is, an OU process with an extremely weak attraction parameter) to force capture of a different speed of movement to that of our actual BM state. It therefore became necessary to include an additional BM state to better model that diversity of movement and allow the subordinate states to represent the social behaviours. This baboon analysis presented a similar scenario.

The MCMC algorithm was initialised as in Section 4.1. λ_{\max} was again set as 0.2, which seems sufficient given the resulting transition rates. Updates of the animal's trajectory are performed over lengths of 3–40 observations and we sampled 210 trajectories in each iteration. The MCMC ran for 500,000 iterations, of which 100,000 was burn-in, and we recorded every second iteration for the movement parameters and every 20th iteration for the behaviour parameters. An uninformative Dirichlet prior was used for the transition rates.

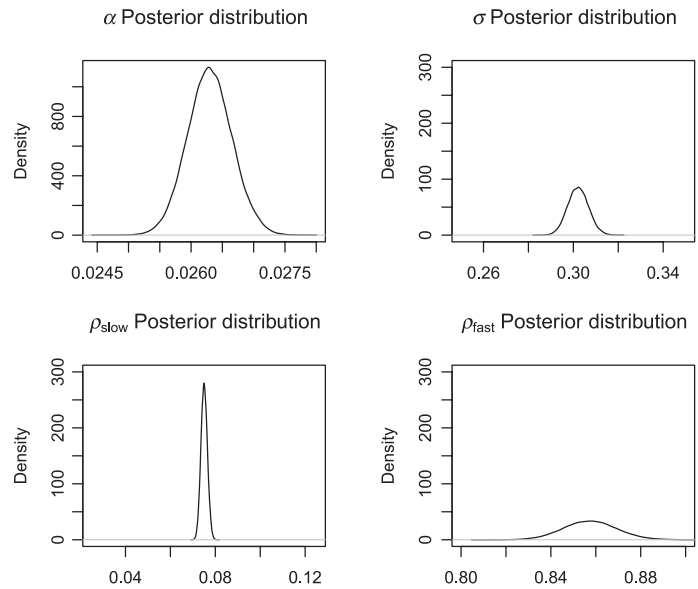
Figure 4 shows the posterior distributions for the movement parameters. The posterior for α shows strong evidence that there is indeed interaction between the five baboons in the data analysed. The Gelman-Rubin diagnostic is again used to check convergence over two MCMC runs (using *coda*, Plummer et al., 2006) and the multivariate potential scale reduction factor of the movement parameters is 1.02.

The state posteriors of baboon 5 are shown in Figure 5, through which we can observe the dynamics of this baboon's social behaviour in a subordinate sense. That is, it mostly alternates being attracted to baboons 3 and 4, though there is some uncertainty as to which baboon it is subordinate to in the last 200 observations.

Figure 7a shows a similar graph but with the state posteriors reconfigured to indicate the role of this particular baboon (baboon 9). An animal in a BM state is *leading* if they have a subordinate or *independent* if not. All animals in a subordinate state are *following*. This allows us to see in what capacity an animal interacts with its peers. For instance, we estimate baboon 9 largely interacts as a subordinate until around observation 700 when it takes on a consistent leadership role. Looking at the data, this time corresponds to a change in the direction that the baboons are moving in.

This formulation also allows us to see the long-term manner in which the animals interact, both in the sense of their role (Table 1) and with each other (Table 2). Though both of these tables only show us a static overview of the social interaction, they highlight which animals

FIGURE 4 Top: posteriors distributions for the four movement parameters for the baboon data analysis. Note the different scales of the *Density* axes. Bottom: a summary of the movement parameter inference. Point estimates and standard deviations are given to 3 s.f.; effective sample size is rounded down after being calculated using *coda* (Plummer et al., 2006)



Parameter	Point estimate	Standard deviation	Effective sample size
α	0.0263	0.000359	1,033
σ	0.3020	0.004620	1,479
ρ_{slow}	0.0751	0.001430	1,620
ρ_{fast}	0.8580	0.011900	5,625

FIGURE 5 The state posterior distribution for baboon 5. There are seven states: the two Brownian motion speeds and five subordinate behaviours where state S_{B_i} indicates attraction to baboon i . The area of each box represents the posterior probability of being in that state at that observation, from 0 to 1




TABLE 1 The percentage (1 d.p.) of observations each baboon spent in each role based on the modal state

	Independent (%)	Leading (%)	Following (%)
Baboon 3	47.2	37.2	15.7
Baboon 4	54.1	10.0	35.9
Baboon 5	20.4	38.0	41.6
Baboon 11	28.7	4.0	67.3
Baboon 9	47.7	21.4	30.9

seem to have persistently high levels of influence in the group (baboons 3 and 5) and potentially strong bonds within it (e.g. baboons 3, 5 and 11).

Exact inference by means of simulating when behaviour switches occur is a computationally costly task. Ideally, we would have run a greater number of iterations as indicated by the modest effective sample size of the movement parameters (Figure 4) and the Gelman-Rubin diagnostic for the transition rates. Whilst the potential scale reduction factor for the vast majority of rates was <1.04 , the highest was 1.38. Additionally, the larger the subset of data we can fit

TABLE 2 The percentage (1 d.p.) of observations each baboon is subordinate to another based on the modal state. Cell ij in the table corresponds to baboon i being subordinate baboon j

	Baboon 3 (%)	Baboon 4 (%)	Baboon 5 (%)	Baboon 11 (%)	Baboon 9 (%)
Baboon 3	—	0.2	7.8	0.1	7.2
Baboon 4	7.8	—	18.4	6.5	3.0
Baboon 5	28.5	10.0	—	1.3	1.9
Baboon 11	16.1	0.0	34.6	—	15.9
Baboon 9	13.9	0.0	9.8	7.2	—

the model to (both in terms of the number of animals and the time period), the more long-term and biological questions can be investigated. However, the results described above were obtained over approximately 2.8 days and so we have not yet overcome the computational hurdle that can come with continuous-time models, particular ones with a complex state space.

This data has previously been analysed by Strandburg-Peshkin et al. (2015) to investigate how a troop of baboons collectively make their movement decisions. ‘Movement initiations’ were extracted from the data of 25 baboons through a method based on minima and maxima distances between a pair of animals. Through this approach, Strandburg-Peshkin et al. quantified the probability of a baboon following a movement initiation in the context of the number of initiators and their consensus in direction—situation-based covariates that would not be trivial to implement in our model.

Though, whilst their interaction assumption of a following animal moving towards the initiator is conceptually similar to ours, it is perhaps more restrictive. Our OU approach models the subordinate’s movement as being distributed around the location of its dominant. As a result, we do not have to constrain our definition of influence to being the ‘initiator’ or the animal at the front of the group (Strandburg-Peshkin et al., 2015) for example. Along with the state switching we have defined, this will also help to smooth out erroneous interactions in the data.

5 | RELIABILITY

Whilst the above results are encouraging, there is scope for the model to infer false positives (interaction where there is none) and false negatives (no interaction where there is some). To investigate this, we have analysed 400 different simulations derived from four different sets of parameters—100 from each set.

The parameters of set 1 are identical to those used to create the simulated data in Section 4.1; set 2 has been derived to simulate similar movement behaviours as inferred from the baboon data; set 3’s parameters were chosen to represent different behaviours from those of sets 1 and 2. For instance, the three noise parameters are less distinct than in the other sets to provide a tougher inference scenario (see Table 3). Set 4 consists of animals solely moving in Brownian motion in order to analyse the rate at which the model

TABLE 3 The four sets of movement parameter values used to simulate data

Parameter set	α	σ	ρ_{slow}	ρ_{fast}
1	0.5	0.7	0.4	1.8
2	0.0266	0.297	0.0751	0.855
3	0.2	0.7	0.5	1.4
4			0.4	1.8

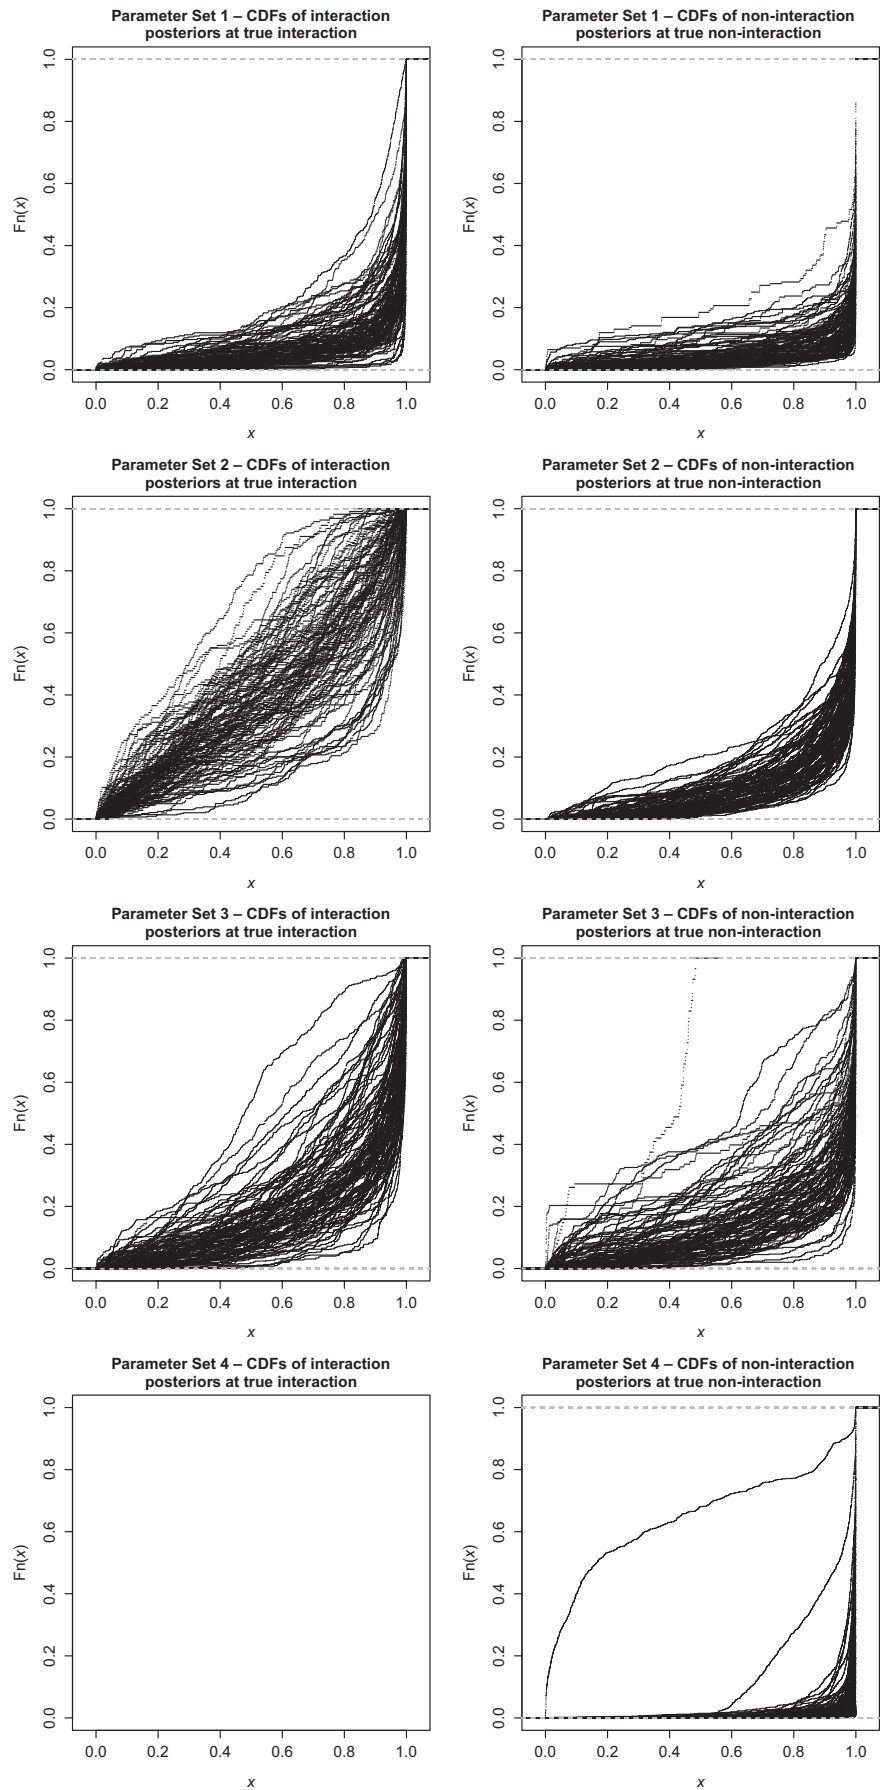
introduces false positives. Whilst the movement parameters and Λ were kept constant across all simulations for that parameter set, the initial behaviour states at time 0 were sampled from the stationary distribution of Λ in order for each data set to contain different behaviours. All simulations allow for two BM states and consist of 100 movement steps.

Generally, each run was initialised much like those in Section 4.1. However, each set was treated to its own movement parameter proposal distributions in order to encourage good mixing, with the same configuration being used for all runs corresponding to that set. Additionally, to encourage false positives in set 4, we initialised α to be low on the assumption that any false positives inferred would point to weak attraction. Each run was carried out for 1 million iterations with a burn-in of 500,000.

In order to evaluate the rate of false positives and false negatives in interaction, we use our state estimations of the whole group to calculate the probability of two animals interacting at a given time, whether directly or indirectly and regardless of dominance and subordination ordering. This probability is an *interaction posterior*. At times of true interaction in the simulations, we would expect the *interaction posterior* to be close to 1; at times of true non-interaction, we would expect the complimentary *non-interaction posterior* to be close to 1. The CDFs of these posteriors for each run are plotted in Figure 6.

The CDFs relating to sets 1 and 3 form the desired curve, though those of set 3 display more uncertainty. This is intuitive as α is smaller in set 3 compared to set 1 (that is, the attraction is weaker) and the noise parameters were set to provide more of a challenge. The anomalous result in the *non-interaction posteriors* for set 3 comes from a data set where there is very little non-interaction—each pair of animals don’t interact for only approximately 8.7% of the data on average. Uncertain state estimations during some of those segments of non-interaction are a blot on otherwise reasonable results for that data set. With regards to set 4, there are no plots concerning true interaction as there isn’t any in those simulations—a feature confidently estimated in the *non-interaction posteriors*. However, there are clearly two erroneous results. The most spurious of those is the consequence of not tuning each simulation separately and the resulting acceptance rate of the movement parameters was extremely low (approximately 1%). Whilst this is clearly not a desirable outcome, ordinarily we would be able to tune the MCMC differently to navigate this issue; as it is, the problem is evident from the most cursory diagnostics, so would

FIGURE 6 The plots in the left column contain the CDFs of the *interaction posteriors* at times of true interaction. The plots in the right column contain the CDFs of the *non-interaction posteriors* at times of true non-interaction. Each row of plots corresponds to a parameter set and each CDF is derived from the posteriors for all pair combinations during a single run



not mislead in an actual analysis. The second erroneous result simply inferred some interaction. A small number of segments of the data are estimated, with a large degree of uncertainty, to contain weak interaction.

Figure 6 indicates a greater degree of uncertainty in the *interaction posteriors* of set 2 compared to those of sets 1 and 3. In part, this effect will be due to set 2's simulations containing comparatively little interaction and so false negatives are more pronounced in their *interaction posterior* CDFs than those of sets 1 and 3. More importantly though, this uncertainty highlights a limitation of modelling the movement of leading/independent animals only as Brownian motion. The difference in certainty of the state posteriors in Section 4.2 and those of the simulations derived from Section 4.2's results will in part be due to there being some feature of the real movement that is not captured by BM, such as persistence. Without that feature, independent animals and subgroups in the simulations would frequently overlap movement paths. It therefore becomes a challenge to identify the exact interactions taking place. This produces scenarios in the simulations that are perhaps not biologically relevant and another process, such as Ornstein-Uhlenbeck, may be better suited to model the movement of the leading and independent animals.

Generally though, these simulation results display a good level of robustness. Where there is uncertainty, the CDFs in Figure 6 indicate it is most likely regarding interaction—potentially resulting in a false negative. Using set 2 as an example, we can examine this uncertainty further with Figures E.1 and E.2 in Appendix E. The *role posteriors* compared against the true role show that we are fairly certain whether an animal is a subordinate or not at any given time. The consistent movement parameter posteriors are further evidence of this. The uncertainty relates to exactly which animal they are subordinate to.

6 | DATA THINNING

Alongside our model being slow to fit to data, a recurring question in animal tracking studies is: at what frequency should data be collected at for the analysis at hand (Hughey, Hein, Strandburg-Peshkin, & Jensen, 2018)? We have therefore experimented with thinning the same baboon data analysed above (henceforth referenced as the 'full analysis') to investigate how the inference speed can be improved by fitting the model to a thinned data set, and how the results compare to the full analysis. We experimented by thinning the data by a factor of five and of 20.

We will also discuss our choices of λ_{\max} for these experiments. $\lambda_{\max}\delta t$ is the mean number of potential switching points sampled from the Poisson process between sequential observations over a time span of δt . Whilst we state in Section 3.1 that $\lambda_{\max} \geq \max(\lambda_{ij})$ to ensure all state switches can be sampled appropriately, if $\lambda_{\max}\delta t$ is high, say five, there is little information to be gained from sampling that many state switches between sequential observations. Reducing λ_{\max} will reduce the size of the augmented data set and the number of computations needed.

6.1 | Thinning by a factor of five

We took every fifth observation of the baboon data detailed above, leaving us with 180 observations (at 0.2 Hz) instead of 899 (at 1 Hz). We ran the inference again for 500 K iterations to provide a comparable analysis. λ_{\max} was kept at 0.2 and so $\lambda_{\max}\delta t = 1$ for sequential observations.

As a result of thinning of the data, the uncertainty of movement parameter posteriors increased as can be seen in Figure F.1 in Appendix F. The slight differences in the estimations will most likely be for two reasons. Firstly, any discrepancies in state estimations between the two analyses and secondly, the model we have chosen to fit may not be sufficient to represent some of the movement behaviours. That is, we chose to model the movement of leading/independent animals as Brownian motion. However, if there is some persistence in their movement in reality, the noise coefficient of BM will not scale appropriately as the data is thinned. Section 2.3 detailed how this model may be extended to capture that persistence.

When comparing the role posteriors of baboon 9, the thinned analysis (Figure 7b) smooths over the state estimations from the full analysis (Figure 7a). To investigate this smoothing, we looked at the modal state estimations at the 1 Hz observation times in both analyses and quantified the periods of behaviour that were present in the full but absent from the thinned analysis. To classify a period of behaviour as being smoothed over, we require that at least 80% of that period to have a different modal state estimated in the thinned analysis. The 80% threshold is arbitrary, but it prevents a period from being accounted for in the thinned analysis by a relatively short visit to that state whilst still capturing significant smoothing. Figure 8 shows that the smoothed-over periods are predominately small in length—mostly under five seconds which is our new temporal scale.

This smoothing is likely to be a result of our assumption that the animal's behaviour is Markovian. If the full data suggests the short state transitions as seen in Figures 5 and 7a, the Markovian behaviour process will readily accommodate them. The constant transition rates incorporate no presumption against very short visits, as reflected in the exponentially distributed holding times. Through thinning the data, the 'evidence' of these transitions is no longer present, though they may remain with much lower probability.

6.2 | Thinning by a factor of 20

We took every 20th observation of the data set used for the full analysis, leaving us with 45 observations (at 0.05 Hz). We ran the inference again for 500 K iterations but for this analysis we reduced λ_{\max} to 0.05. As discussed above, the motivation for this was to keep $\lambda_{\max}\delta t$ small.

Whilst this amount of thinning can be thought of as extreme smoothing, where we'd expect to get results largely in agreement to those of the full analysis, Figure 7c shows the state/role posteriors can be drastically different. For example, the thinned data

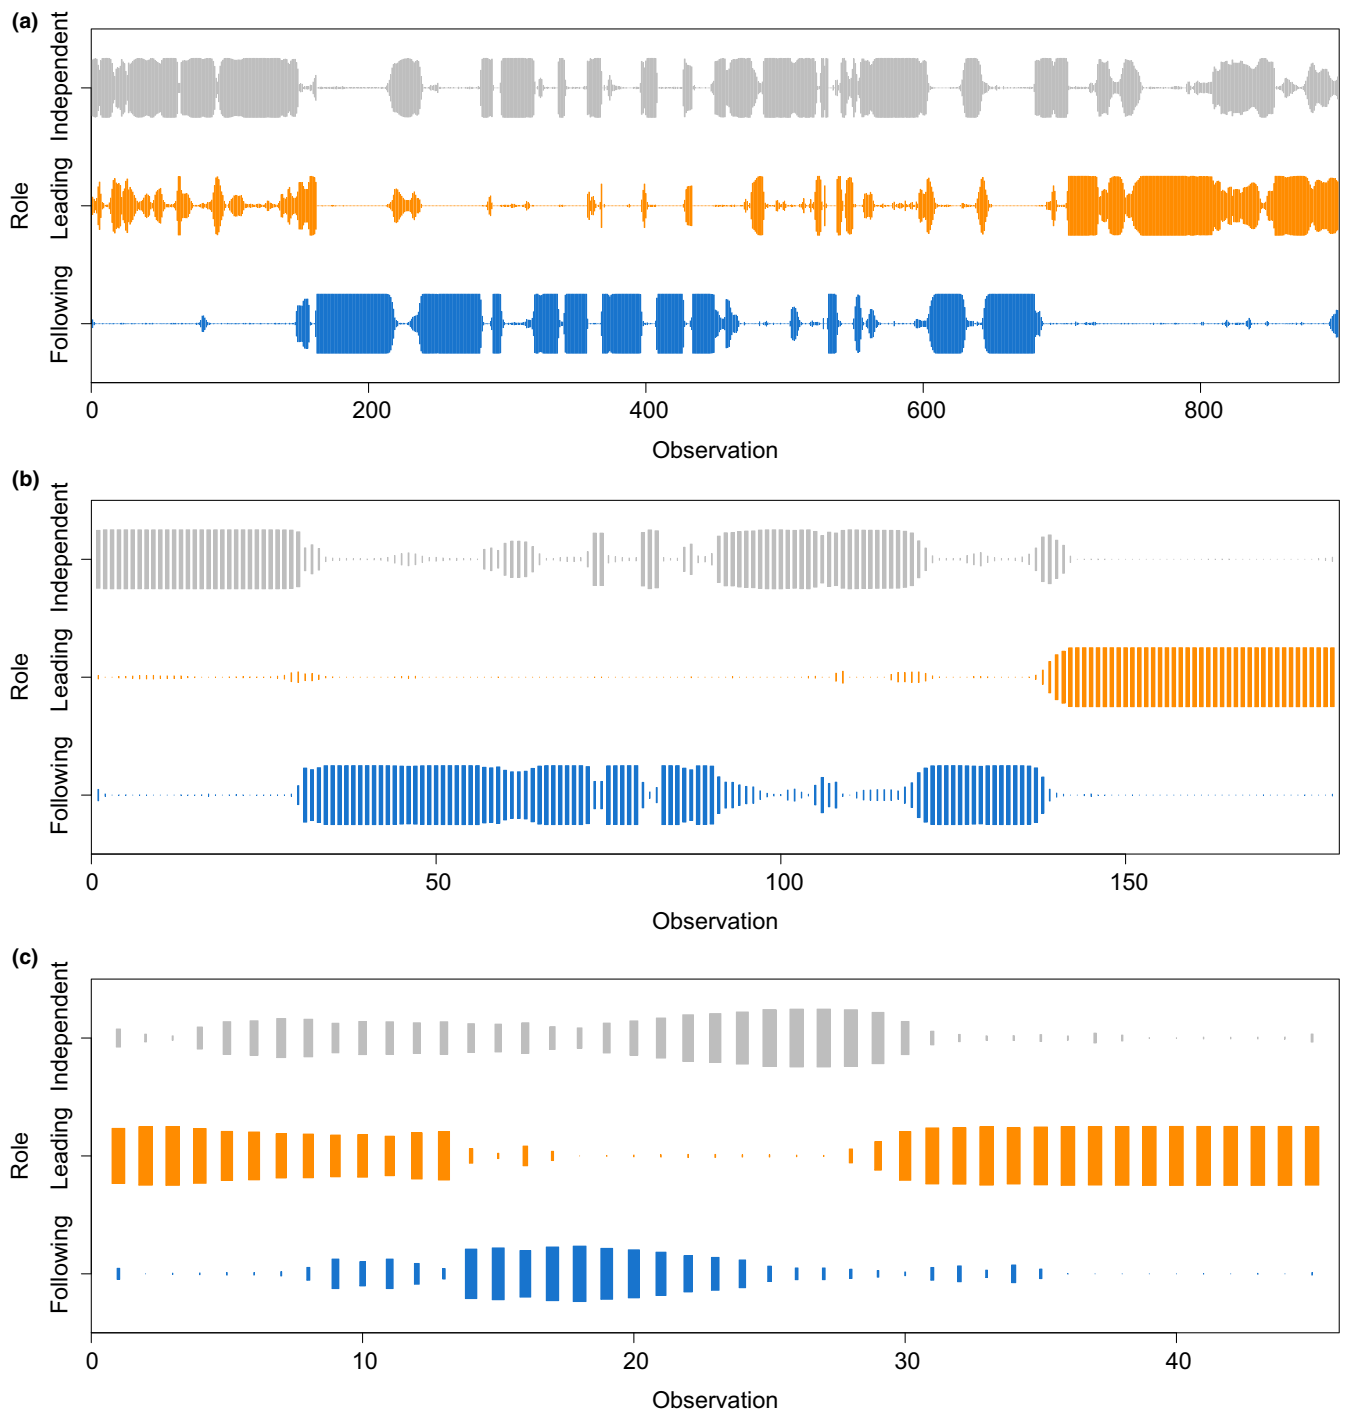


FIGURE 7 The role posterior distribution for baboon 9 for the full analysis (a), thinned by a factor of five (b) and thinned by a factor of 20 (c). The area of each box represents the posterior probability of being in that role at that observation, from 0 to 1

no longer supports the two periods of subordination that baboon 9 undertakes between observations 600–700 in Figure 7a. Thus, the resulting social structure has a different picture. These different behaviour estimations then have further impacts on the movement parameter posterior distributions (Figure F.1, Appendix F).

In the full analysis, $\max(\lambda_u) = 0.153$ (3 d.p) and so we wouldn't expect an animal to switch state much more frequently than every 7 s. We therefore suspect thinning the data to 0.2 Hz keeps enough information in order to estimate similar behaviours.

Thinning the data further, and reducing λ_{\max} , reduces the scope of social behaviours that can be captured. That is, short-term interactions can't be seen and are unlikely to be sampled. Whilst the thinning-by-20 results aren't wrong per se, they offer a different temporal perspective of the social behaviours of the baboons. These thinning experiments indicate the results of an analysis will be most informative and efficient if the temporal resolution of the data collected is informed by the nature of the behaviours that are to be investigated.




FIGURE 8 A histogram of the number of observations in a state before switching in the full analysis (black) and which of these periods were smoothed over in the thinned analysis (grey)

With regards to the speed of the inference, the thinned-by-five analysis was completed in approximately 24 hr, as opposed to approximately 67 hr for the full analysis. Whilst the number of observations was reduced to a fifth, the number of simulated data was equivalent as we kept the same λ_{\max} . Though, there is scope to use a smaller λ_{\max} during the thinned-by-five analysis as the resulting $\max(\lambda_{ij})$ was 0.111 (3 s.f.). The thinned-by-20 analysis was completed in approximately 7 hr.

7 | COMPARISON OF METHODS

To compare our approach with other methods, we have examined our results from both the simulation and full baboon analyses alongside those obtained when applying dyadic metrics (Joo, Etienne, Bez, & Mahévas, 2018; Long et al., 2014) to the same data. Much like our approach, dyadic metrics utilise movement data from multiple animals to investigate any interdependence and in turn better understand their collective behaviours. As their name suggests, the dyadic metrics are built to analyse the interdependence of two animals.

The metrics we have chosen to implement are *proximity* and *dynamic interaction* using the *wildlifeDI* R package from Long et al. (2014). Joo et al. (2018) indicate each metric offers an informative view into a specific element of interaction. *Proximity* evaluates whether two animals are within a distance threshold defined by the user; *dynamic interaction*, when split into its *displacement* and *direction* components, offers insight into whether two animals are moving at a similar speed or in a similar orientation respectively. In order to compare results, we use the same *interaction posterior* approach that we used in Section 5 as that is comparable to the interpretation of the dyadic metrics. This is done for every data point in which both animals are simultaneously observed as this is a necessary criterion for the metrics.

Figure 9a plots our *interaction posterior* against each of the three metrics for animals 1 and 3. These animals were chosen as ideal candidates to compare methods as they undertook periods

of both direct and indirect interaction. Whilst the *interaction posterior* is largely (and correctly) concentrated at 0 and 1, the *displacement* and *direction* metrics are uniform across their ranges regardless of whether there was true interaction or not. *Proximity* fares better as most simultaneous data <2 m apart correspond to true interaction and most simultaneous data >4 m apart correspond to true non-interaction. However, the distances in between highlight the difficulty in determining interaction from *proximity* alone.

Figure 9b displays the same comparison for baboons 4 and 9. Whilst the *interaction posterior* is again largely concentrated towards 0 and 1, the *proximity* metric does not offer a discernible pattern. It is uninformative when the *interaction posterior* is close to 0 and interaction appears to be grouped at distinct ranges of proximity—at approximately 60 m and <20 m. For both *displacement* and *direction*, there is moderate concentration at (1,1)—indicating a consensus between our model and these metrics on some moments of interaction. However, both metrics are quite uniform when the *interaction posterior* is certain there is no interaction. In particular, *direction* is just as likely to suggest interaction as it is non-interaction.

Overall, there is some consistency of the metrics with our model in estimating when two animals interact. However, when the animals are not interacting (either known from the simulated data or estimated from our model), the dyadic metrics are generally not in agreement and are fairly uniform across their ranges. We suspect this is because the metrics can have quite relaxed definitions of interaction. For instance, *dynamic interaction* can hint at interaction without any concern as to whether the animals are reasonably proximate to one another. That is in contrast to our perhaps more defined notion of interaction where we are not likely to introduce false positives (see Section 5). Though, that definition comes with its own limitations such as not being able to capture co-movement.

From a practical point of view, the dyadic metrics are much faster to run than our approach and they will require less tuning. However, our model can more naturally handle incomplete or unsynchronised data. The metrics require essentially simultaneous data to evaluate the cohesion of two animals at a given

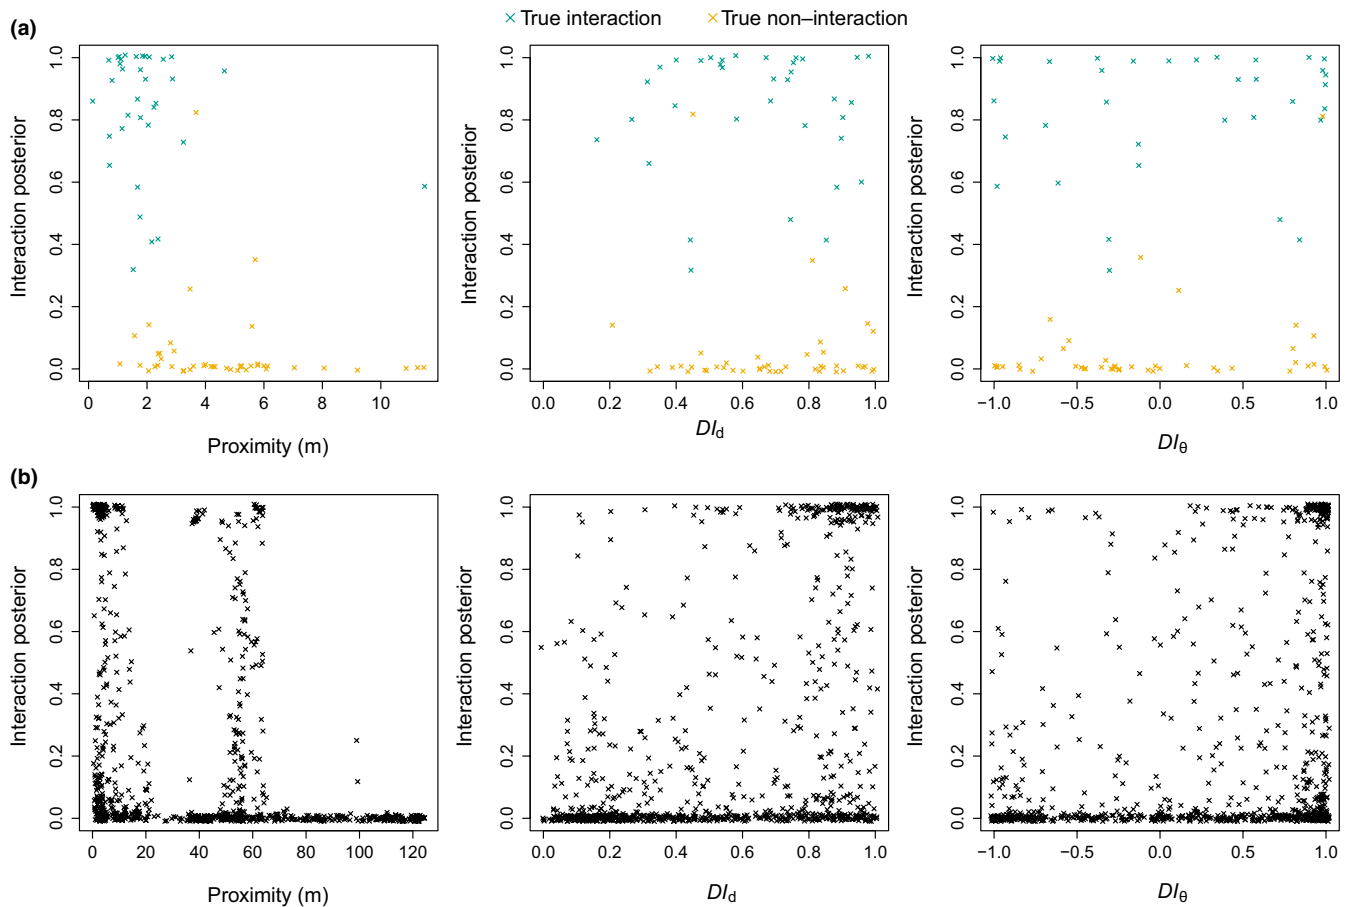


FIGURE 9 The *interaction posteriors* (the posterior probability of two animals being in the same subgroup) from the simulation analysis (Section 4.1) are plotted against the *proximity*, *dynamic interaction in displacement* (DI_d) and *dynamic interaction in direction* (DI_θ) dyadic metrics for the same data set (a). The results used are for animals 1 and 3 and each point corresponds to a simultaneous observation. Similarly, the same approach has been taken for baboons 4 and 9 from Section 4.2 (b). Both axes have been jittered in order to help display the density of the points

time whereas our model is able provide an estimate of interaction in continuous time across the temporal range of the data. Furthermore, our model has the capability to jointly analyse the data of larger social groups.

8 | DISCUSSION

Our model offers a robust yet flexible method in continuous time in which to capture the social interactions in animal movement. Modelling the state switches in continuous time means we avoid approximation from using discrete data and our social behaviour analysis is not bound to the temporal scale of the data. Through the ‘influence hierarchies’ framework, a wide range of social constructs can be captured: from despotic leadership to fission-fusion dynamics. Furthermore, social hierarchies are a simple and common concept meaning a certain amount of the abstraction of statistical models is stripped away.

Our formulation of the behavioural states allows us to capture rich information with regards to the social behaviours in a group’s movement. Animals consistently influencing the movement of other

animals may be considered ‘keystone’ animals and identifying such animals can have productive applications in conservation and management decisions (King et al., 2018; Westley et al., 2018). Obtaining a picture of a group’s social structure may also help us understand how resilient or adaptive they are to change (King et al., 2018) and we can monitor how anthropogenic activity might be impacting them (Westley et al., 2018).

We restrict the social structure to the hierarchies defined in Section 2.1 in order to keep the model tractable. However, this does mean we omit certain interactive behaviours. For instance, the ‘double-subordinate’ as in McDonald and Shizuka (2013) and discussed in Strandburg-Peshkin, Papageorgiou, Crofoot, and Farine (2018), where an animal is equally influenced by two (or more) others. Cyclic structures are also not accounted for in our model, but analysis from McDonald and Shizuka (2013) suggests social structures tend to be highly orderly and so this may not be problematic.

Our social interaction assumption is one based on order—that is, there is a dominant and subordinate. However, that will not always be a fair assumption such as the case where two or more animals are mutually influenced by each other. Current work is ongoing to

model that behaviour in a similar framework to the one detailed in this paper. Furthermore, our assumption the behaviour process in general is Markovian may not be biologically sound as discussed in Section 6.1. An alternative approach may be to model the behaviour process as semi-Markovian in order for the times spent in a state to be more realistically distributed.

Whilst the opportunity to obtain simultaneous tracking data from multiple animals within a group is increasing, it still may not be possible or practical to track an entire group. Incomplete data on a group, along with our definitions of the behaviour states, can lead to falsely estimating direct influence and relationships between the animals we have data on as our model does not account for the animals we don't have data for. For example, we have characterised baboon 9 as leading others through a movement decision, but as we have only analysed five of the 26 baboons in the group there is a possibility baboon 9 was in fact influenced by an animal we don't have tracking data for. Therefore we have to be conscious not to over-interpret our results as the de-facto social behaviour. Niu et al. (2016) used a moving abstract central point that acts as the focus of a group's movement which can navigate the possibility that we don't have complete data about a group, but it doesn't allow finer details of the social behaviour to be examined.

Whilst probing further into the social interactions can offer richer information as discussed, the sheer size of the state space can be inhibiting—a considerable amount of time is needed to fully explore it. Work needs to be done to optimise the current inference algorithm, particularly with regards to the method in which we sample new trajectories. The current method is naive as we simulate forwards through the data whilst ignoring the end state the trajectory must be in, which can lead to low acceptance rates. One potential route to remedy this is to explore techniques based on the Forward Algorithm (widely used for efficient implementation of Hidden Markov models) as proposed by Blackwell (2018).

Various elements of the model can be treated as either heterogeneous or homogeneous in space and time or across individuals, depending on the analysis in question. In both analyses in this paper, we used a Λ that is heterogeneous across individuals. An alternative approach would be to treat it as homogeneous, resulting in just four transition rate parameters relating to being in an OU (subordinate state) or a BM (leading/independent) state. This should improve the mixing of the inference but it will provide coarser results. The model and inference we have detailed is spatially homogeneous. However, more rich and beneficial information could be obtained through adding spatial covariates such as environmental data or the positions of the animals relative to each other. These may affect the social behaviour, or independent and/or leading animals may be affected by a static (Blackwell et al., 2016) or dynamic (Wang, Blackwell, Merkle, & Potts, 2019) environment rather than following Brownian motion. The algorithm in Section 2.3 to sample new trajectories is easily extended to a spatially heterogeneous case, though it slows it down considerably. Again, work is ongoing to develop and improve this. Furthermore, we could add some heterogeneity into the animals

themselves. For example, we could use characteristic information such as the animal's age or sex to ascertain how they interplay with social dominance. Or independent transition rates could be estimated for each animal in a model that is hierarchical in the statistical sense, but some consideration will be needed to weigh up the benefits of substantially increasing the number of parameters.

ACKNOWLEDGEMENTS

We are grateful to Eloise Bray for discussion about dyadic metrics and to several referees for comments on earlier versions of this paper. J.E.M. is supported by a studentship from the Engineering and Physical Sciences Research Council.

AUTHORS' CONTRIBUTIONS

J.E.M. conceived the model; J.E.M., P.G.B. and M.N. developed the methodology and implementation; J.E.M. led the analysis, with contributions from P.G.B. J.E.M. led the writing; all authors contributed critically to the writing and gave final approval for publication.

DATA AVAILABILITY STATEMENT

The R code required to run the inferences in Sections 4, 5 and 6, along with the simulated data, is available at the following GitHub repository <https://github.com/jemilner/influenceHierarchy> (Milner, Blackwell, & Niu, 2020, <https://doi.org/10.5281/zenodo.3972134>). The baboon data used in Section 4.2 is available on Movebank baboon. The particular subset that we used relates to baboons 3, 4, 5, 9 and 11 between the times 05:51:38 and 06:06:36 on 2012-08-03 (YYYY-MM-DD).

ORCID

Jordan E. Milner  <https://orcid.org/0000-0002-0863-3158>

Paul G. Blackwell  <https://orcid.org/0000-0002-3141-4914>

Mu Niu  <https://orcid.org/0000-0002-3068-5501>

REFERENCES

- Bivand, R. S., Pebesma, E., & Gomez-Rubio, V. (2013). *Applied spatial data analysis with R* (2nd ed.). New York, NY: Springer.
- Blackwell, P. G. (1997). Random diffusion models for animal movement. *Ecological Modelling*, 100, 87–102. [https://doi.org/10.1016/S0304-3800\(97\)00153-1](https://doi.org/10.1016/S0304-3800(97)00153-1)
- Blackwell, P. G. (2003). Bayesian inference for Markov processes with diffusion and discrete components. *Biometrika*, 90, 613–627. <https://doi.org/10.1093/biomet/90.3.613>
- Blackwell, P. G. (2018). *Integrated continuous-time hidden Markov Models*. Retrieved from <https://arxiv.org/abs/1807.11907>
- Blackwell, P. G., Niu, M., Lambert, M. S., & LaPoint, S. D. (2016). Exact Bayesian inference for animal movement in continuous time. *Methods in Ecology and Evolution*, 7, 184–195. <https://doi.org/10.1111/2041-210X.12460>
- Chase, I. D., & Lindquist, W. B. (2016). The fragility of individual-based explanations of social hierarchies: A test using animal pecking orders. *PLoS ONE*, 11, 1–16. <https://doi.org/10.1371/journal.pone.0158900>
- Crofoot, M., Kays, R., & Wikelski, M. (2015). Data from: Shared decision-making drives collective movement in wild baboons, <https://doi.org/10.5441/001/1.kn0816jn>

- Dunn, J. E., & Gipson, P. S. (1977). Analysis of radio telemetry data in studies of home range. *Biometrics*, 33, 85–101. <https://doi.org/10.2307/2529305>
- Harris, K. J., & Blackwell, P. G. (2013). Flexible continuous-time modeling for heterogeneous animal movement. *Ecological Modelling*, 255, 29–37. <https://doi.org/10.1016/j.ecolmodel.2013.01.020>
- Hughey, L. F., Hein, A. M., Strandburg-Peshkin, A., & Jensen, F. H. (2018). Challenges and solutions for studying collective animal behaviour in the wild. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373, 20170005. <https://doi.org/10.1098/rstb.2017.0005>
- Johnson, D. S., London, J. M., Lea, M.-A., & Durban, J. W. (2008). Continuous-time correlated random walk model for animal telemetry data. *Ecology*, 89, 1208–1215. <https://doi.org/10.1890/07-1032.1>
- Joo, R., Etienne, M.-P., Bez, N., & Mahévas, S. (2018). Metrics for describing dyadic movement: A review. *Movement Ecology*, 6, 26. <https://doi.org/10.1186/s40462-018-0144-2>
- Kays, R., Crofoot, M. C., Jetz, W., & Wikelski, M. (2015). Terrestrial animal tracking as an eye on life and planet. *Science*, 348, <https://doi.org/10.1126/science.aaa2478>
- King, A. J., Fehlmann, G., Biro, D., Ward, A. J., & Fürtbauer, I. (2018). Re-wilding collective behaviour: An ecological perspective. *Trends in Ecology & Evolution*, 33, 347–357. <https://doi.org/10.1016/j.tree.2018.03.004>
- Langrock, R., Hopcraft, J. G. C., Blackwell, P. G., Goodall, V., King, R., Niu, M., ... Schick, R. S. (2014). Modelling group dynamic animal movement. *Methods in Ecology and Evolution*, 5, 190–199. <https://doi.org/10.1111/2041-210X.12155>
- Langrock, R., King, R., Matthiopoulos, J., Thomas, L., Fortin, D., & Morales, J. M. (2012). Flexible and practical modeling of animal telemetry data: Hidden Markov models and extensions. *Ecology*, 93, 2336–2342. <https://doi.org/10.1890/11-2241.1>
- Long, J. A., Nelson, T. A., Webb, S. L., & Gee, K. L. (2014). A critical examination of indices of dynamic interaction for wildlife telemetry studies. *Journal of Animal Ecology*, 83, 1216–1233. <https://doi.org/10.1111/1365-2656.12198>
- Martin, A. D., Quinn, K. M., & Park, J. H. (2011). MCMCpack: Markov chain Monte Carlo in R. *Journal of Statistical Software*, 42, 22.
- McDonald, D. B., & Shizuka, D. (2013). Comparative transitive and temporal orderliness in dominance networks. *Behavioral Ecology*, 24, 511–520. <https://doi.org/10.1093/beheco/ars192>
- Milner, J. E., Blackwell, P. G., & Niu, M. (2020). jemilner/influenceHierarchy: Code for MEE article: Modelling and inference for the movement of interacting animals. <https://doi.org/10.5281/zenodo.3972134>
- Niu, M., Blackwell, P. G., & Skarin, A. (2016). Modeling interdependent animal movement in continuous time. *Biometrics*, 72, 315–324. <https://doi.org/10.1111/biom.12454>
- Pebesma, E. J., & Bivand, R. S. (2005). Classes and methods for spatial data in R. *R News*, 5, 9–13.
- Plummer, M., Best, N., Cowles, K., & Vines, K. (2006). Coda: Convergence diagnosis and output analysis for MCMC. *R News*, 6, 7–11.
- R Core Team. (2017). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Ramos-Fernández, G., & Morales, J. M. (2014). Unraveling fission–fusion dynamics: How subgroup properties and dyadic interactions influence individual decisions. *Behavioral Ecology and Sociobiology*, 68, 1225–1235. <https://doi.org/10.1007/s00265-014-1733-8>
- Strandburg-Peshkin, A., Farine, D. R., Couzin, I. D., & Crofoot, M. C. (2015). Shared decision-making drives collective movement in wild baboons. *Science*, 348, 1358–1361. <https://doi.org/10.1126/science.aaa5099>
- Strandburg-Peshkin, A., Papageorgiou, D., Crofoot, M. C., & Farine, D. R. (2018). Inferring influence and leadership in moving animal groups. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373, 20170006. <https://doi.org/10.1098/rstb.2017.0006>
- Wang, Y.-S., Blackwell, P. G., Merkle, J. A., & Potts, J. R. (2019). Continuous time resource selection analysis for moving animals. *Methods in Ecology and Evolution*, 10, 1664–1678. <https://doi.org/10.1111/2041-210X.13259>
- Westley, P. A., Berdahl, A. M., Torney, C. J., & Biro, D. (2018). Collective movement in ecology: From emerging technologies to conservation and management. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373. <https://doi.org/10.1098/rstb.2017.0004>

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section.

How to cite this article: Milner JE, Blackwell PG, Niu M. Modelling and inference for the movement of interacting animals. *Methods Ecol Evol* 2021;12:54–69. <https://doi.org/10.1111/2041-210X.13468>